



# UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE  
United States Patent and Trademark Office  
Address: COMMISSIONER FOR PATENTS  
P.O. Box 1450  
Alexandria, Virginia 22313-1450  
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/858,080	05/15/2001	Kevin Collins	10006721-1	2538

  

EXAMINER	
BATURAY, ALICIA	

  

ART UNIT	PAPER NUMBER
2155	

  

MAIL DATE	DELIVERY MODE
11/16/2007	PAPER

7590 11/16/2007  
HEWLETT-PACKARD COMPANY  
Intellectual Property Administration  
P.O. Box 272400  
Fort Collins, CO 80527-2400

**Please find below and/or attached an Office communication concerning this application or proceeding.**

The time period for reply, if any, is set in the attached communication.

40



UNITED STATES PATENT AND TRADEMARK OFFICE

Commissioner for Patents  
United States Patent and Trademark Office  
P.O. Box 1450  
Alexandria, VA 22313-1450  
[www.uspto.gov](http://www.uspto.gov)

**MAILED**

**NOV 16 2007**

**Technology Center 2100**

**BEFORE THE BOARD OF PATENT APPEALS  
AND INTERFERENCES**

Application Number: 09/858,080

Filing Date: May 15, 2001

Appellant(s): COLLINS ET AL.

\_\_\_\_\_  
Philip S. Lyren  
(Registration # 40,709)  
For Appellant

**EXAMINER'S ANSWER**

This is in response to the appeal brief filed 08 August 2007 appealing from the Office action  
mailed 21 March 2007.

**(1) Real Party in Interest**

A statement identifying by name the real party in interest is contained in the brief.

**(2) Related Appeals and Interferences**

The examiner is not aware of any related appeals, interferences, or judicial proceedings which will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

**(3) Status of Claims**

The statement of the status of claims contained in the brief is correct.

**(4) Status of Amendments After Final**

The appellant's statement of the status of amendments after final rejection contained in the brief is correct.

**(5) Summary of Claimed Subject Matter**

The summary of claimed subject matter contained in the brief is correct.

**(6) Grounds of Rejection to be Reviewed on Appeal**

The appellant's statement of the grounds of rejection to be reviewed on appeal is correct.

Art Unit: 2155

**(7) Claims Appendix**

The copy of the appealed claims contained in the Appendix to the brief is correct.

**(8) Evidence Relied Upon**

U.S. 6,157,963

Courtright, II et al.

12-2000

**(9) Grounds of Rejection**

The following ground(s) of rejection are applicable to the appealed claims:

Claims 1-21 are rejected under 35 U.S.C. 102(e) as being anticipated by Courtright, II et al. (U.S. 6,157,963).

**DETAILED ACTION**

***Claim Rejections - 35 USC § 102***

1. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(e) the invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent, except that an international application filed under the treaty defined in section 351(a) shall have the effects for purposes of this subsection of an application filed in the United States only if the international application designated the United States and was published under Article 21(2) of such treaty in the English language.

2. Claims 1-21 are rejected under 35 U.S.C. 102(e) as being anticipated by Courtright, II et al. (U.S. 6,157,963).

3. With respect to claim 1, Courtright teaches a method for managing transactions at a network storage device comprising:

Receiving an incoming transaction at the network storage device (Courtright, col. 4, lines 52-58); and assigning a priority to the incoming transaction relative to other incoming transactions at the network storage device based at least in part on a usage policy (Courtright, col. 4, lines 58-67); and overriding the priority with a requested priority included in the incoming transaction (Courtright, col. 8, lines 24-42).

4. With respect to claim 2, Courtright teaches the invention described in claim 1, including a method further comprising receiving the usage policy at a network storage device (Courtright, col. 8, lines 31-42), and where the network storage device is a NAS device (Courtright, col. 1, lines 13-19).

5. With respect to claim 3, Courtright teaches the invention described in claim 1, including a method further comprising:

Reading meta data from the transaction; and comparing the meta data to a number of rules defined in the usage policy, where assigning the priority to the transaction is based on at least part of the meta data satisfying at least one condition of the number of rules (Courtright, col. 4, lines 52-58).

6. With respect to claim 4, Courtright teaches the invention described in claim 1, including ordering the transaction among other transactions in a queue at the network storage device (Courtright, col. 8, line 63 – col. 9, line 3).
7. With respect to claim 5, Courtright teaches a method for managing transactions at plural network storage devices, comprising:  
  
Generating a usage policy at a server for the network storage devices (Courtright, col. 8, lines 31-36); distributing the usage policy from the server across a network to the network storage devices for prioritizing a plurality of incoming transactions received at the network storage devices relative to one another; and providing updates to the usage policy from the server to the network storage devices (Courtright, col. 8, lines 36-42).
8. With respect to claim 6, Courtright teaches the invention described in claim 5, including a method further comprising identifying the network storage devices on the network (Courtright, col. 1, lines 13-19), and where the network storage devices are NAS devices (Courtright, col. 1, lines 13-19).
9. With respect to claim 7, Courtright teaches the invention described in claim 5, including a method where the usage policy comprises a number of rules, each including meta data and a corresponding priority (Courtright, col. 4, lines 52-67).

10. With respect to claim 8, Courtright teaches an apparatus for managing a plurality of incoming transactions at a network storage device, comprising:

Computer readable storage medium at the network storage device; a usage policy stored on the computer readable medium; and computer readable program code residing in the computer readable storage medium, comprising program code for prioritizing the plurality of incoming transactions relative to one another based on the usage policy (Courtright, col. 4, lines 52-67); where the prioritizing in the usage policy uses at least two conditions based on (1) user logon, (2) originating application, (3) user-requested priority, and (4) purpose for accessing the network storage device (Courtright, col. 8, lines 24-30).

11. With respect to claim 12, Courtright teaches the invention described in claim 8, including an apparatus where the usage policy comprises a number of default rules (Courtright, col. 6, line 1 – col. 7, line 3).

12. With respect to claim 13, Courtright teaches an apparatus for managing a plurality of incoming and outgoing transactions at a network storage device, comprising:

Computer readable storage medium (Courtright, col. 3, lines 55-61); and computer readable program code residing in the storage medium, including program code for defining a usage policy for prioritizing the plurality of incoming and outgoing transactions relative to one another (Courtright, col. 5, lines 7-30).

Art Unit: 2155

13. With respect to claim 14, Courtright teaches the invention described in claim 13, where the computer readable program code resides at a policy management server and further comprises program code for distributing the usage policy to the network storage device (Courtright, col. 8, lines 31-42).

14. Claims 9-11, 15, 16 and 18-21 do not teach or define any new limitations above claims 1-3 and 13 and therefore are rejected for similar reasons.

#### **(10) Response to Argument**

The examiner summarizes the various points raised by the appellant and addresses them individually.

*(A) Appellant Argues:* Nowhere does Courtright teach or even suggest overriding priorities with a requested priority included in an incoming transaction. Courtright never states that this user assigned priority overrides a previous priority that was assigned at the network storage device.

*In Response:* The examiner respectfully submits that Courtright teaches overriding priorities with a requested priority included in an incoming transaction (the user of the system may change the priority schemes and/or algorithms used by the system, and the user may change the priorities assigned to the storage users and/or storage objects...dynamically during I/O



Art Unit: 2155

processing...the system is dynamically configurable – see Courtright, col. 8, lines 24-42). Because a user can change the priorities assigned *during* I/O processing, this new user-defined priority can be seen as overriding and thus taking the place of the previous priority present on the network storage device.

**(B) *Appellant Argues:*** Claim 1 recites that the priority is assigned “at” the network storage device. By contrast, Courtright assigns the priority at the storage controller, not the storage object or network storage device.

***In Response:*** The examiner respectfully submits that Courtright teaches the storage controller also acting as a storage object. “Moreover, while the present invention is described herein as processing each I/O request directly with a storage object, one skilled in the art will appreciate that caching, write-back, and other time and system resource saving techniques may be utilized. For example, if one or more blocks of data are accessed frequently by users, that data may be maintained in a cache memory and accessed directly from there. This type of configuration will prevent a large number of processor intensive I/O operations to the same storage object. Similarly, instead of holding a storage user until the storage controller completes a write operation, the processor may use write-back techniques that release the storage user once the write data is received from the storage user” (see Courtright, col. 5, lines 55-67). Thus the storage controller can respond both to requests for reading data out of its cache, and requests for writing data by using write-back techniques that temporarily store user-made changes to the data

Art Unit: 2155

on the storage controller. Thus the priority assignments are made at the storage controller, which functions both as a storage controller and as a storage object, and therefore priority is assigned at a network storage device.

(C) ***Appellant Argues:*** The claim then recites “distributing said usage policy from said server across a network to said network storage devices.” This element is not taught in Courtright. Courtright never discloses that the prioritizing algorithm is generated at a server and then distributed to the network storage devices as recited in claim 5. In Courtright, the priorities are never distributed to the storage devices. Looking to Figure 2, Courtright is teaching that a user can change the priorities and download these changes to the storage controller 16.

***In Response:*** The examiner respectfully submits that in the aforementioned paragraph (B), it is established that the storage controller also functions as a storage object, and that priority is therefore assigned at a network storage device. Additionally, Courtright teaches the prioritizing algorithm is generated at a server (the user can change the storage user and/or storage object priorities manually via a user interface) and then distributed to the network storage devices (the priority values may be downloaded to the system and in particular to the storage controller...through a network environment – see Courtright, col. 8, lines 36-42). Thus the priorities are distributed to the storage controller, which also functions as a network storage device, via a network.

**(D) Appellant Argues:** Claim 5 recites “providing updates to said usage policy from said server to said network storage devices.” Thus, Courtright never teaches that updates are provided from the server to the storage devices. In Courtright, the updates are provided to the server, not the storage devices.

**In Response:** The examiner respectfully submits that in the aforementioned paragraph (B), it is established that the storage controller also functions as a storage object, and that priority is therefore assigned at a network storage device. Additionally, Courtright teaches providing updates to said usage policy (the user can change the storage user and/or storage object priorities manually via a user interface) from said server to said network storage devices (the priority values may be downloaded to the system and in particular to the storage controller...through a network environment – see Courtright, col. 8, lines 36-42). Thus the updates are distributed to the storage controller, which also functions as a network storage device, via a network.

**(E) Appellant Argues:** Claim 8 recites that the prioritizing in the usage policy “uses at least two conditions based on (1) user logon, (2) originating application, (3) user-requested priority, and (4) purpose for accessing the network storage device.” Nowhere does Courtright teach that priority is based on at least two conditions as recited in claim 8.

**In Response:** The examiner respectfully submits that Courtright teaches that “a plurality of prioritizing schemes may be used to prioritize and schedule the I/O requests. I/O requests from

Art Unit: 2155

storage users are first prioritized by storage user and then by storage object. Processor receives the I/O requests and determines, using a first prioritization factor which is storage user priority, in which one of the memory queues in the first storage area it should place the I/O request. After placing one or more I/O requests in the particular memory queues in the first storage area, processor then retrieves I/O requests from the memory queues and determines, using a second prioritization factor which is storage object priority, in which one of the memory queues in a second storage area it should place the I/O requests (see Courtright, col. 8, lines 44 – col. 9, line 3).” Because Courtright outlines this method of utilizing a plurality of prioritizing schemes in which I/O requests are prioritized by one factor and then another, this demonstrates the “prioritizing in the usage policy uses at least two conditions” limitation of the claim. Courtright also states that “I/O requests may be prioritized by storage user priority, storage object priority, job priority or a combination of all of these factors (see Courtright, col. 8, lines 24-30),” of which storage user priority is equivalent to the (1) user login limitation and job priority is equivalent to the (4) purpose for accessing the network storage device limitation in the claim. Therefore Courtright can be shown to address both the specific types of priority, and the two conditions of priority called for in the claim.

**(F) Appellant Argues:** Courtright teaches priority based on incoming I/O requests. Courtright never mentions how priority is determined for outgoing I/O requests. Courtright never discloses a relationship for prioritizing both incoming and outgoing I/O requests “relative to one another.”

***In Response:*** The examiner respectfully submits that Courtright teaches "I/O requests from storage users are first prioritized by storage user and then by storage object. Processor receives the I/O requests and determines, using a first prioritization factor which is storage user priority, in which one of the memory queues in the first storage area it should place the I/O request. After placing one or more I/O requests in the particular memory queues in the first storage area, processor then retrieves I/O requests from the memory queues and determines, using a second prioritization factor which is storage object priority, in which one of the memory queues in a second storage area it should place the I/O requests (see Courtright, col. 8, lines 44 – col. 9, line 3)." This shows incoming transactions being prioritized relative to one another. After this process is completed, "processor retrieves an I/O request from memory queues in second storage area in accordance with a predetermined scheduling algorithm, and processes the I/O operation for the selected I/O request with the appropriate storage object(s). The processor then returns the results of the I/O operation back to the requesting client (see Courtright, col. 9, lines 12-18). Thus, as transactions are prioritized coming into the system, the results returned are to the client in the order in which they are processed, which is in accordance with a scheduling algorithm. Therefore by virtue of the transactions being prioritized as they come into the system, the order in which they are subsequently sent out is determined.

Art Unit: 2155

**(11) Related Proceeding(s) Appendix**

No decision rendered by a court or the Board is identified by the examiner in the Related Appeals and Interferences section of this examiner's answer.

For the above reasons, it is believed that the rejections should be sustained.

Respectfully submitted,



Alicia Baturay

7 November 2007

Conferees:

/Lynne H Browne/  
Lynne H. Browne  
Appeal Practice Specialist, TQAS  
Technology Center 2100

Saleh Najjar  
Supervisory Patent Examiner  
Art Unit 2155



**SALEH NAJJAR**  
**SUPERVISORY PATENT EXAMINER**